Supplementary Information: Neural-network solutions to stochastic reaction networks

Ying Tang,[1,*] Jiayu Weng,[2,†] and Pan Zhang[3,4,5,‡]

[1] *International Academic Center of Complex Systems, Beijing Normal University, Zhuhai 519087, China*
[2] *Faculty of Arts and Sciences, Beijing Normal University, Zhuhai 519087, China*
[3] *CAS Key Laboratory for Theoretical Physics, Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing 100190, China*
[4] *School of Fundamental Physics and Mathematical Sciences, Hangzhou Institute for Advanced Study, UCAS, Hangzhou 310024, China*
[5] *International Centre for Theoretical Physics Asia-Pacific, Beijing/Hangzhou, China*

In the Supplementary Information, we provide the application of the method to more examples in addition to those in the main text, including the birth-death process, gene expression without regulation, an autoregulatory feedback loop, and more cases of the intracellular signaling cascade. We also list the computational details of all the models in the table.

CONTENTS

I. Birth-death process 1

II. Gene expression 2

III. Autoregulatory feedback loop 3

IV. More cases of the intracellular signaling cascade 4

V. Table of the computational details 5

References 5

I. BIRTH-DEATH PROCESS

As the first example, we consider the time-homogeneous birth-death process [1] with the count of the species in the interval $X \in [0, N]$. The reactions are $\emptyset \xrightarrow{k_2} X$, and $X \xrightarrow{k_1} \emptyset$. Its dynamics are described by a time-dependent probability distribution $P_t(n)$ with the count n at time t. The chemical master equation is:

$$\partial_t P_t(n) = B(n+1)P_t(n+1) + F(n-1)P_t(n-1) - [B(n) + F(n)]P_t(n), \quad (0 \le n \le N), \tag{S1.1}$$

where $B(n)$ and $F(n)$ are backward (death) and forward (birth) propensities, respectively. We consider the case of propensities: $F(n) = k_2$, $B(n) = k_1 n$ with k_1, k_2 as rates, for $0 \le n \le N$ and zero otherwise, with the boundary conditions $F(N) = 0$, $B(0) = 0$.

The birth-death process has an analytical solution. With the Poisson initial distribution:

$$P_0(n) = \frac{e^{-\alpha_0}\alpha_0^n}{n!}, \tag{S1.2}$$

where α_0 is a parameter, the time-dependent distribution is [1]:

$$P_t(n) = \frac{e^{-\alpha_t}\alpha_t^n}{n!}, \tag{S1.3}$$

$$\alpha_t = \alpha_0 e^{-k_1 t} + (k_2/k_1)(1 - e^{-k_1 t}). \tag{S1.4}$$

* These authors contributed equally; Corresponding authors: jamestang23@gmail.com
† These authors contributed equally
‡ Corresponding authors: panzhang@itp.ac.cn

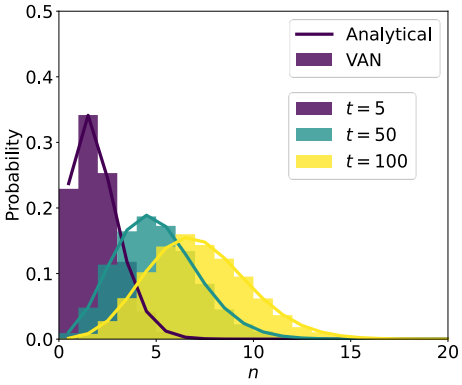

Supplementary Figure 1. The result from the VAN for the birth-death process. The distributions of the species between the VAN and the analytical solution match over time. The bar is the result from the VAN, and the line is the analytical solution Eq. (S1.3). The color specifies the time points $t = 5, 50, 100$, and the transparency is used to visualize the overlap. Parameters: $k_1 = 0.01$, $k_2 = 0.1$. The hyperparameters of the VAN are in Supplementary TABLE II, and the initial distribution is the Poisson distribution.

The analytical solution makes it convenient to test the present numerical approach. In Supplementary Fig. 1, the distributions from the analytical solution and the present approach match at several time points. This shows that the present method can track the time evolution of the probability distribution for the birth-death process.

II. GENE EXPRESSION

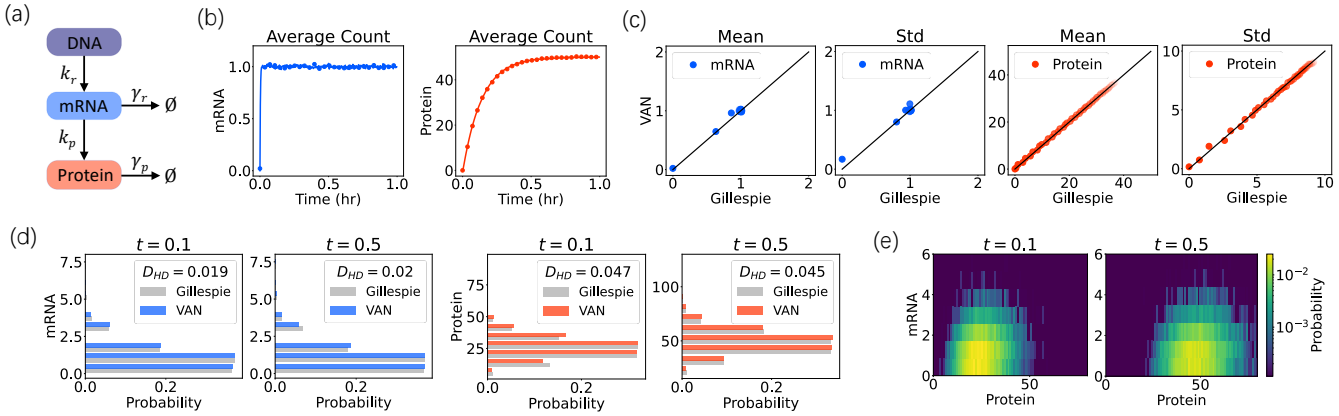

Supplementary Figure 2. The result of gene expression without regulation. (a) A schematic of the chemical reaction. (b) The time series of the average counts for the mRNA and protein specified by color, from the VAN (dots) and the Gillespie simulation (lines). (c) Comparison of the means and standard deviations of the mRNA and protein between the VAN and the Gillespie simulation, at time points $t = 0, 0.004, 0.008, ..., 0.196, 0.2$ denoted by the transparency. (d) The marginal distributions of the mRNA and protein at time points $t = 0.1, 0.5$, from the Gillespie simulation (gray) and the VAN with the same color in (b). The inset shows the Hellinger distance between the two distributions. (e) The joint distribution of mRNAs and proteins at time points $t = 0.1, 0.5$ from the VAN, with the color as the probability values in the logarithmic scale. The parameters are $k_r = 0.1$, $k_p = 0.1$, $\gamma_r = 0.1$, and $\gamma_p = 0.002$, and the Gillespie simulation has 10^4 trajectories. The hyperparameters of the VAN are in Supplementary TABLE II, and the initial distribution is the delta distribution with zero mRNA and protein.

We consider the example of gene expression without regulation, as one of the cornerstone examples in the systems biology. The system involves two species, mRNA and protein. The four reactions include the transcription from DNA to mRNA, the transcription from mRNA to protein, and the decay of the mRNA and protein. The parameters are the transcription rate k_r, transcription rate k_p, and decay rates γ_r, γ_p.

This example of gene expression has the chemical reaction:

$$DNA \xrightarrow{k_r} r, \quad r \xrightarrow{k_p} r + p, \quad r \xrightarrow{\gamma_r} \emptyset, \quad p \xrightarrow{\gamma_p} \emptyset. \tag{S2.1}$$

where r denotes the mRNA, p denotes the protein, and $k_r, k_p, \gamma_r, \gamma_p$ are rate constants. The parameter values used in the simulation are listed in Supplementary Fig. 2.

In Supplementary Fig. 2, we show the mean, the marginal distribution and the joint distribution of the counts of the mRNA and protein over time, by using Gillespie simulation or the VAN method. The marginal statistics match well between the two methods. We quantify the similarity of the two distributions by the Hellinger distance, which is below 0.05. This indicates that the two distributions are similar, validating the accuracy of the VAN in generating the distribution.

III. AUTOREGULATORY FEEDBACK LOOP

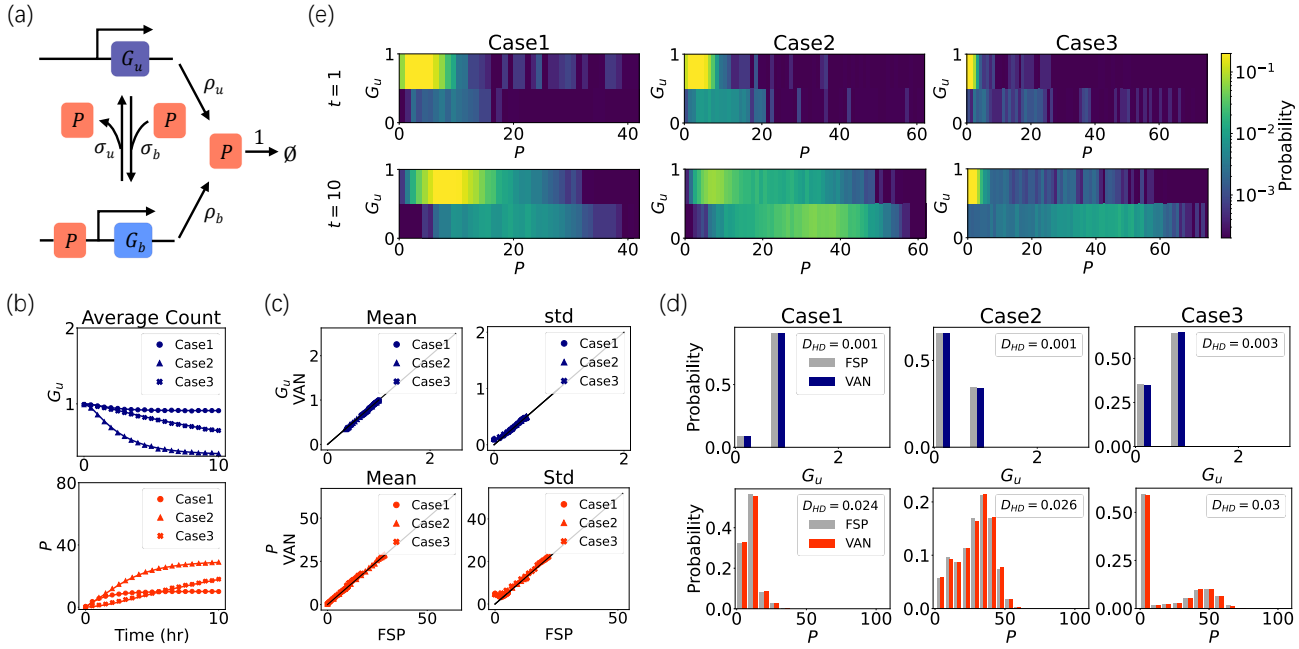

Supplementary Figure 3. The result of the autoregulatory feedback loop. (a) A schematic of the reactions. (b) The time evolution of the average counts for the promoter (blue) and protein (red), from the VAN (symbols) and the finite state projection (FSP) (lines). Three cases with different parameters in Supplementary TABLE I are marked with different symbols (circle, triangle and cross). (c) Comparison on the means and standard deviations of the promoter and protein between the VAN and the FSP, at time points $t = 0, 0.5, ..., 9.5, 10$. (d) The marginal distributions of the promoter and protein at time points $t = 10$ from the FSP (gray) and the VAN with the same color in (b). The inset shows the Hellinger distance between the two distributions. (e) The joint distribution of the promoter and protein calculated from the VAN at time points $t = 1, 10$ in different cases, with the color bar for the probability values in the logarithmic scale. The values of the parameters in the three cases are shown in Supplementary TABLE I. The hyperparameters of the VAN are in Supplementary TABLE II, and the initial distribution is the delta distribution with one G_u and zero protein.

We consider an autoregulatory feedback loop [2] as illustrated in Supplementary Fig. 3a. The gene has two promoter states G_b, G_u switching to each other, with a binding rate σ_b and an unbinding rate σ_u. The two states of bond and unbound have different translation rates. The model can be written as the chemical reactions:

$$G_u + P \xrightarrow{\sigma_b} G_b, \quad G_b \xrightarrow{\sigma_u} G_u + P, \quad G_u \xrightarrow{\rho_u} P, \quad G_b \xrightarrow{\rho_b} P, \quad P \xrightarrow{1} \emptyset. \tag{S3.1}$$

where $\sigma_b, \sigma_u, \rho_b, \rho_u$ are rate constants. The total count of the two promoter states is conserved: $G_b + G_u = 1$. This conservation effectively reduces a variable by having $G_u = 1 - G_b$, and sets a constraint on the counts of the two promoters: $G_u = 0, 1$, $G_b = 0, 1$. We have put this constraint on the neural network of the VAN. The three groups of parameters under the consideration [2] are listed in Supplementary TABLE I.

Parameter	σ_u	σ_b	ρ_u	ρ_b
Case 1	0.94	0.01	8.40	28.1
Case 2	0.69	0.07	7.20	40.6
Case 3	0.44	0.08	0.94	53.1

Supplementary Table I. The values of the parameters used in the three cases of the autoregulatory feedback loop.

By the present approach, the means (Supplementary Fig. 3b), the standard deviations (Supplementary Fig. 3c) and the marginal distributions (Supplementary Fig. 3d) of the counts of the protein and promoter match those from the finite state projection, for the three sets of parameters. As shown in the marginal distribution (Supplementary Fig. 3d) and the joint distribution (Supplementary Fig. 3e), the three cases include unimodal and bimodal distributions for the count of the protein. The results demonstrate that the VAN generates accurate marginal statistics and can effectively produce the bimodal probability distribution for the reaction network with the feedback loop.

IV. MORE CASES OF THE INTRACELLULAR SIGNALING CASCADE

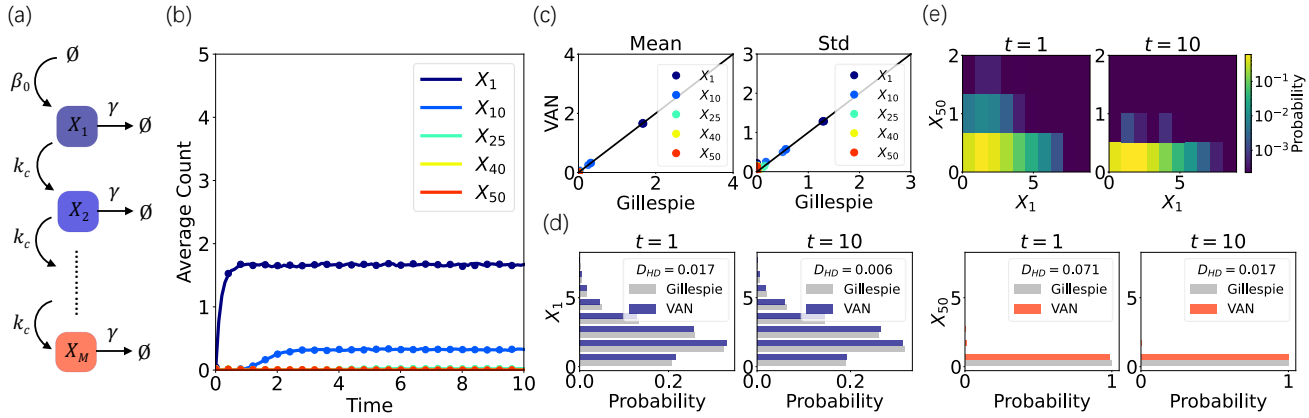

Supplementary Figure 4. The result of the linear signaling cascade (case 1) with 50 species. (a) A schematic of the chemical reaction. (b) The time series of the average count of species from the VAN (dots) and the Gillespie simulation (lines). The color specifies the chemical species. (c) Comparison of the means and standard deviations of the chemical species between the VAN and the Gillespie simulation, at time points $t = 1, 2, ..., 9, 10$ for chemical species denoted by color. (d) The marginal count distributions of various chemical species are plotted horizontally at time points $t = 1, 10$. The color in (d) specifies the result from the VAN for two species, and gray denotes the Gillespie method. The inset contains the Hellinger distance between the two distributions. (e) The joint distribution of the first and the last species at time points $t = 1, 10$ from the VAN, where the color bar denotes the probability values in the logarithmic scale.

We provide the details of cases 2, 3 for the intracellular signaling cascade as considered in [3], to complement the case 1 in the main text.

As in the case 2, the signaling cascade with nonlinear activation has the reaction scheme in Eq. (12), and replaces the constant catalytic rate by a Hill function with a basal rate. For each reaction $X_i \longrightarrow X_{i+1}$, the catalytic production is: $H(x_i) = b + (k_m x_i^h)/(k_0 + x_i^h)$, where b is the basal rate, k_m controls the strength of the Hill activation, k_0 corresponds to the affinity for the substrate, and h is the Hill coefficient. For various M, the count of the last species x_M changes dramatically. For a better comparison, we have rescaled the parameters to be M dependent, such that the count of the last species is kept similar for various M. We have also rescaled the time scale by 25-fold, such that the system reaches the steady state with a similar simulation time as the case 1.

In the case 3, the signaling cascade with negative feedback has the same reactions in Eq. (12), and an additional negative feedback from the last species to the first one: $H(x_M) = b + k_m/(k_0 + x_M^h)$, where the parameters have the same meaning of the case 2.

The result of the case 1 with $M = 50$ species is shown in Supplementary Fig. 4. It shows a good match of the marginal statistics between the Gillespie simulation and the VAN method, demonstrating that the present approach is applicable to high-dimensional systems. The results of cases 2, 3 with $M = 15$ species are separately shown in Supplementary Figs. 7, 8. The accuracy of the VAN can be further improved by using more epochs, especially for the case 2 where the count of each species depends sensitively on the precedent species with nonlinear activation. The

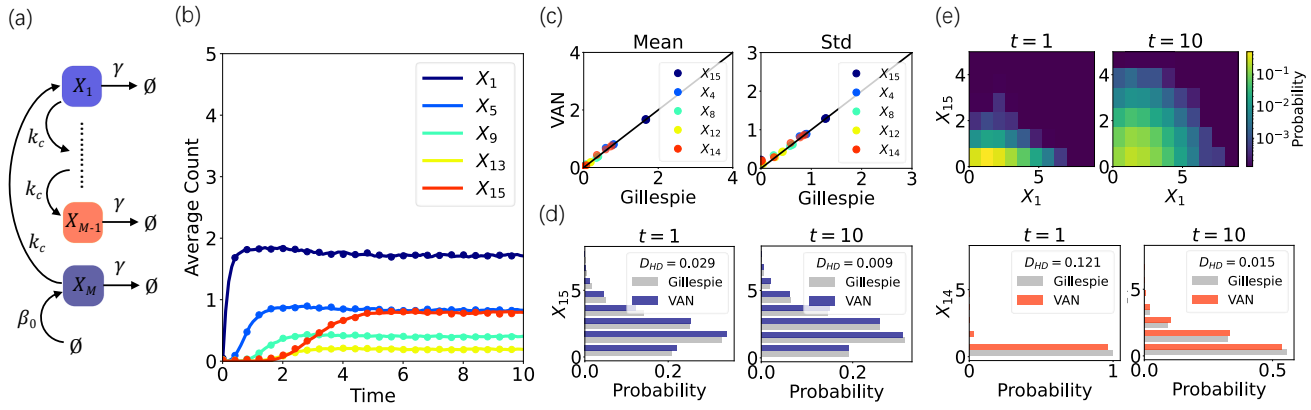

Supplementary Figure 5. The result of the linear signaling cascade (case 1) with the switched order of the 15 species. The signaling starts from the last species X_{15}, activating X_1 and so on. The result is as accurate as that without switching the order. The layout of the figure is the same as that of Supplementary Fig. 4.

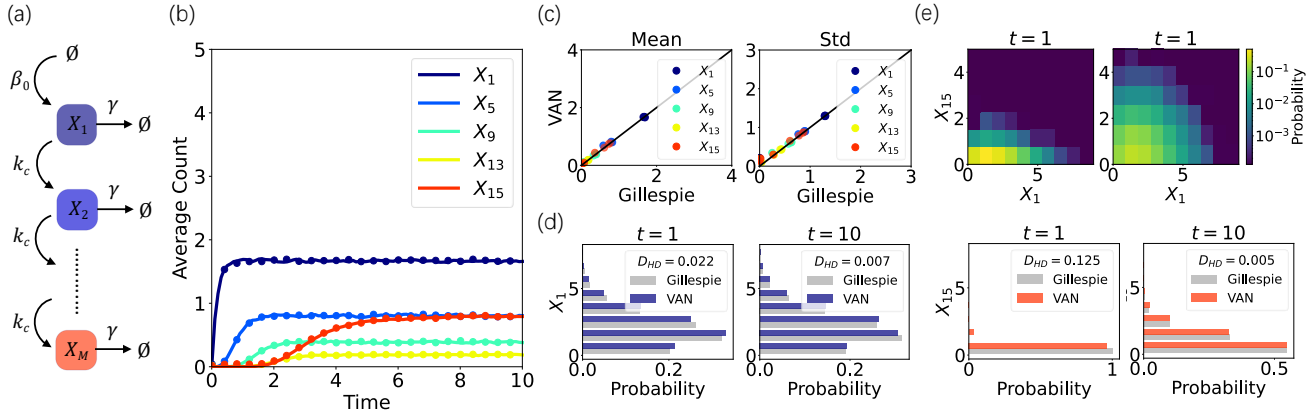

Supplementary Figure 6. The result of the linear signaling cascade (case 1) with 15 species by using the transformer. The layout of the figure is the same as that of Supplementary Fig. 4.

results indicate that the present approach is applicable to a high-dimensional system with nonlinear interactions and feedback.

V. TABLE OF THE COMPUTATIONAL DETAILS

We list the hyperparameters of the VAN and the corresponding computational time for all examples from the RNN in Supplementary TABLE II. The results from the transformer for the example of the signaling cascade are summarized in Supplementary TABLE III. The computational times were recorded in hour scales sublinearly to the number of species and reactions (Fig. 5).

[1] C. W. Gardiner, *Handbook of Stochastic Methods*, 3rd ed. (Springer-Verlag, Berlin, 2004).

[2] A. Sukys, K. Öcal, and R. Grima, Approximating solutions of the chemical master equation using neural networks, iScience **25**, 105010 (2022).

[3] A. Gupta, C. Schwab, and M. Khammash, Deepcme: A deep learning framework for computing solution statistics of the chemical master equation, PLoS Comput. Biol. **17**, e1009623 (2021).

[4] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980 (2014).

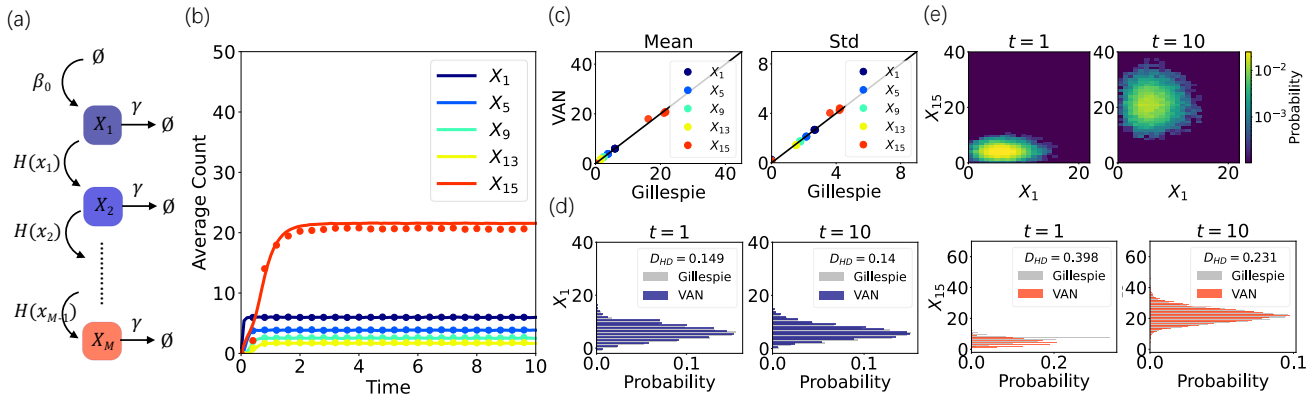

Supplementary Figure 7. The result of the nonlinear signaling cascade (case 2) for 15 species. The captions of each figure are the same as those of Supplementary Fig. 4. The parameters $\beta_0 = 10$, $\gamma = 0.1$, $b = 1$, $k_m = 100$, $k_0 = 10$, $h = 1$. The upper count limit is chosen as $N = 40$, and the result with $N = 100$ has the same accuracy.

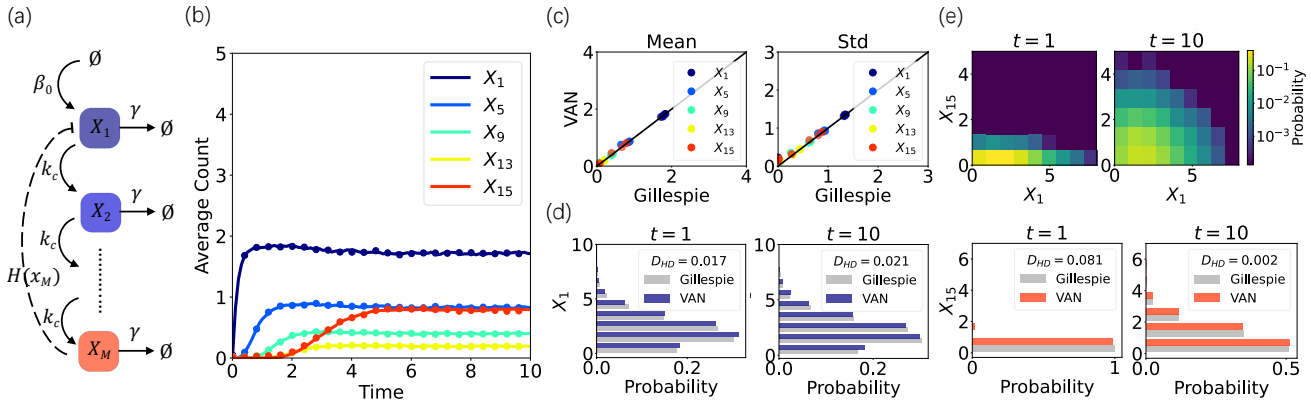

Supplementary Figure 8. The result of the linear signaling cascade with feedback (case 3) for 15 species. The captions of each figure are the same as those of Supplementary Fig. 4. The parameters $\beta_0 = 1$, $k_c = 5$, $\gamma = 1$, $b = 1$, $k_m = 100$, $k_0 = 10$, $h = 1$ [3].

	Species	Reactions	N	Time steps	dt	Epochs	Depth	Width	Comput. time (hr)
Birth-death process	1	2	30	10^4	10^{-2}	100	1	16	1.62
Gene expression	2	4	100	10^4	10^{-1}	100	1	16	3.18
Autoregulatory feedback	2	5	100	$2 * 10^3$	$5 * 10^{-3}$	100	1	32	0.65
Toggle switch	4	8	80	$8 * 10^3$	$5 * 10^{-3}$	100	1	32	4.21
Early life self-replicator	3	6	100	$2 * 10^5$	$4 * 10^{-4}$	20	1	32	42.21
Epidemic model	3	6	80	$1 * 10^4$	$2 * 10^{-2}$	100	1	8	3.68
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	1	32	1.66
Signaling cascade 2	15	30	40	$2 * 10^3$	$5 * 10^{-3}$	100	1	32	3.42
Signaling cascade 3	15	31	10	$1 * 10^3$	10^{-2}	100	1	32	1.66

Supplementary Table II. The computational time of various models, under the chosen parameters of the RNN. The time step length is dt, in the unit of the inverse of the reaction rates. The upper count limit of each species is N. The number of species for the signaling cascade is $M = 15$, and the computational time for various numbers of species ($M = 5, 10, 15, 20, 30, 40, 50$) are shown in Fig. 5. The signaling cascade 2 has more time steps, as it requires smaller dt than that of the signaling cascade 1. The critical parameter value for the early life self-replicator here is $\alpha_c = 1$. The epochs is the number of epochs used at each time step. The depth and width are for the RNN. The learning rate is 10^{-3} with the Adam optimizer [4], and the batch size is 1000. All computational are performed with a single core GPU ($\sim 25\%$ usage) of a Tesla-V100.

	Species	Reactions	N	Time steps	dt	Epochs	n_{layers}	n_{heads}	d_{model}	d_{ff}	Comput. time (hr)
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	2	2	16	32	3.96
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	6	2	16	32	10.11
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	2	4	16	32	4.50
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	2	2	32	32	3.97
Signaling cascade 1	15	30	10	$1 * 10^3$	10^{-2}	100	2	2	16	128	4.31

Supplementary Table III. The computational time of the first case of the signaling cascade with $M = 15$ species. A set of hyperparameters of the transformer are considered, and the transformer typically needs longer computational time than the RNN. The other settings of hyperparameters are the same as shown in Supplementary TABLE II.